

OLGA - A DIALOGUE SYSTEM WITH AN ANIMATED TALKING AGENT

Jonas Beskow¹, Kjell Elenius¹ & Scott McGlashan²

¹Department of Speech, Music and Hearing, KTH, Stockholm

²Swedish Institute for Computer Science, Stockholm

now at Ericsson Radio Systems, Stockholm

E-mail: beskow@speech.kth.se, kjell@speech.kth.se, scott.mcglashan@era-t.ericsson.se

ABSTRACT

The object of the Olga project is to develop an interactive 3D animated talking agent. A futuristic application scenario is interactive digital TV, where the Olga agent would guide naive users through the various services available on the network. The current application is a consumer information service for microwave ovens. Olga required the development of a system with components from many different fields: multimodal interfaces, dialogue management, speech recognition, speech synthesis, graphics, animation, facilities for direct manipulation and database handling. To integrate all knowledge sources Olga is implemented with separate modules communicating with a central dialogue interaction manager. In this paper we mainly describe the talking animated agent and the dialogue manager. There is also a short description of the preliminary speech recogniser used in the project.

1. INTRODUCTION

As spoken dialogue systems for simple information services begin to move from the laboratory into the area of technology, research interest is increasing turning to the integration of spoken dialogue interfaces with other modalities such as graphical interfaces. Apart from the general advantages of allowing an alternative input and output modality, speech can compensate for some of the apparent limitations of a graphical interface. Advantages include increased speed of interaction, higher bandwidth (attention and attitude expressed through stress and prosody, etc.) and ability to describe objects not visually present. Conversely, the graphical interface can compensate for limitations of speech e.g. by making immediately visible the effects of actions upon objects and indicating through the display which objects are currently salient for the system.

By including an animated agent in the interface, several positive effects can be anticipated. The system will seem more anthropomorphic, which will make users more comfortable with the dialogue situation. The character can provide a link between the spoken and the visual information domains, being able to refer to graphical items in the interface, using gaze and pointing. Body language, facial expression and gaze can potentially be very useful communication channels in a spoken interface. Furthermore, proper lip-synchronised articulation will improve intelligibility of system utterances, as shown in Beskow et al. [1].

2. THE OLGA PROJECT

In the Olga project, we have developed a multimodal system combining a dialogue interface with a graphical user interface, which provides consumer information about microwave ovens.

2.1 The Domain: Microwave Ovens

An original motivation for the Olga project was to ease the access to electronic information systems for people who are unfamiliar with computers. They still constitute a substantial part of the population in all ages, with a pre-dominance of elderly people. The selected consumer information application indicates the ambition to make Olga an instrument for the general public. Furthermore, the Swedish Consumer Agency (Konsumentverket) was participating during the initial stages of the project, and they provided a database with facts about microwave ovens.

2.2 Four Main Components

The system is composed of four main components: a speech and language understanding component; a direct manipulation interface which provides graphical information and widgets for navigation; an animated talking agent; and a dialogue manager for co-ordinating interpretation and generation in these modalities.

2.3 Previous Research

Compared with previous research in the area, the novelty of Olga lies in that it integrates interactive spoken dialogue, 3-D animated facial expressions, gestures, lip-synchronised audio-visual speech synthesis and a graphical direct manipulation interface. Cassel et al. [2] have modelled speech and gesture in dialogue using two virtual agents, but no user interaction. Katashi & Akikazu [3] employed animated facial expressions, but no gestures, as a back-channelling mechanism in a spoken dialogue system. Thorisson [4] used a 2-D animated character together with input from many sources, including speech and gaze, to model mainly the social aspects of multimodal dialogue interaction. The Waxholm project at the Department of Speech, Music and Hearing [5], which in some aspects can be seen as a predecessor to Olga, uses a human-like face for the talking agent, and utilises, for example, eye-gaze to refer to various on-screen items such as timetables.

The behaviour of the Olga agent is modelled using rules and parameterised templates; for example, the interaction strategies are based on condition-action rules where the

condition part refers the current interactional state as well as user input, and the action part to schematic descriptions of behaviour in language, graphics, and gestures. Similarly, in the animation module, realisation of a particular gesture is achieved by invoking the selected gesture's template and supplying appropriate parameters. This approach has worked well in our task domain allowing the agent's behaviour to be easily and quickly extended, as well as facilitating software maintenance.

3. THE DIALOGUE MANAGER

The dialogue manager is based on techniques developed in a speech dialogue interface for telephone-based information systems in different languages, see McGlashan [6] and Eckert & McGlashan [7].

3.1 A Tri-partite Model

A tri-partite model of interaction is responsible for semantic, task and dialogue interpretation. The semantics component provides a context-dependent interpretation of user input, and is capable of handling anaphora and ellipsis. A task component embodies navigation strategies to efficiently obtain information from the user necessary for successful database access. The dialogue component adopts an 'event-driven' technique for pragmatically interpreting user input, and producing system responses, compare Giachin & McGlashan [8]. On the basis of user input events, it updates a dialogue model composed of system goals and dialogue strategies. The goals determine the behaviour of the system, allowing for confirmation and clarification of user input (to minimise dialogue breakdown), as well as requests for further information (to maximise dialogue progress). The dialogue strategies are dynamic so that the behaviour of the system varies with progress. A more detailed description of the dialogue manager may be found in Beskow & McGlashan [9].

3.2 Multimodality

In order to manage multimodal dialogues, input and output need to be informationally compatible at the dialogue management level. A user may provide input via buttons in the interface and the agent generate a spoken response; or a user may refer linguistically to an object which the agent realised graphically. Consequently, all input and output is represented in the semantic description language used for spoken input, see McGlashan [6]. This language also allows the user to use different modalities in the same response: e.g. clicking on an object, and then speaking a command to apply to it.

3.3 Output Modality Selection

The dialogue manager decides which modality to use for agent output. In general, modality selection is defined in terms of characteristics of the output information, and the expressiveness and efficiency of the alternative modalities for realising it. In practice, selection is determined by rules which specify realisation in the three modalities depending on the action or state which the agent wants to

express. Table 1 provides a simplified representation of rules (the rules can take in account other aspects of the action or state). Goals with a control or feedback function are realised in speech and gesture: for example, success in understanding user input is indicated with a head nodding gesture, while failure is indicated by speaking an explanation of the failure together with raised eyebrows and the mouth turned down. In cases where the database access has required relaxation of product constraints, speech and a 'regret' gesture are realised. Product information itself is presented in speech and graphics: detailed product information is displayed while the agent gives a spoken overview. Finally, a print action is simply indicated with a graphical icon.

Table 1. Output modality selection rules.

| Condition | Speech | Graphics | Gesture |
|-----------------------------|--------|----------|---------|
| reference success state | no | no | yes |
| reference failure state | yes | no | yes |
| constraint relaxation state | yes | no | yes |
| inform action | yes | yes | yes |
| printing action | no | yes | no |

4. THE ANIMATED AGENT

The Olga character is a three dimensional cartoon-like robot lady, that can be animated in real time. It is capable of text-to-speech synthesis with synchronised movements of lips, jaw and tongue. It also supports gesture and facial expression, that can be used to add emphasis to utterances, support dialogue turn-taking, visually refer to other on-screen graphics such as illustrations and tables, and to indicate the system's internal state: listening, understanding, uncertain, thinking (i.e. doing time-consuming operations such as searching a database) etc.

4.1 The Parameterised Polygon Model

The Olga character, see Figure 1, is implemented as a polygon model, consisting of about 2500 polygons, that can be animated at 25 frames per second on a graphics workstation. The character was first created as a static polygon representation of the body and head, including teeth and tongue. This static model was then parameterised, using a general deformation parameterisation scheme. The scheme allows a deformation to be defined by a few basic properties such as transformation type (rotation, scaling or translation), area of influence (a list of vertex-weight pairs that defines which polygon vertices should be affected by the transformation and to what extent) and various control points for normalisation of the deformation. It is then possible to define non-rigid deformations such as jaw opening, lip rounding etc., by combining basic deformations. Not only articulatory parameters, but also control of eyebrows, eyelids and smiling are defined in this manner. The body was parameterised by introduction of rotational joints at the neck, elbows, wrists, fingers etc.




Figure 1. Wireframe and shaded representations of the Olga character.

4.2 Speech and Articulation

One important reason for using an animated agent in a spoken interface is that it actually will contribute, sometimes significantly, to the intelligibility of the speech, given that mouth movements are properly modelled, compare LeGoff et al. [10]. This is especially true if the acoustic environment is bad, due to for example noise or cross-talk, or if speech perception is impeded by hearing impairment. In a recent experiment, we found that the Olga-character increased the overall intelligibility of VCV-stimuli in noise from 30% for synthetic voice only, to 47% for the synthetic voice and synthetic face combination, see Beskow et al. [1].

Articulation is controlled by a rule-based text-to-speech system framework, see Carlson & Granström [11]. Trajectories for the articulatory parameters are calculated using a set of rules that account for co-articulation effects. This rule set was originally developed for an extended version of the Parke model [12], see Beskow [13]. However, the articulation parameters of the Olga character are chosen to conform to those of the extended Parke model. This makes it possible to drive Olga's articulation using the same set of rules. Once the parameter trajectories are calculated, the animation is carried out in synchrony with play-back of the speech waveform, which in turn is generated by a formant filter synthesiser controlled by the same rule-synthesis framework.

4.3 Complex Gestures

Speech movements are calculated on an utterance-by-utterance basis and played back with high control over synchronisation. Body movements and non-speech facial expressions on the other hand, place different requirements on the animation system. Say for example that we want the agent to dynamically change its expression during a user utterance, depending on the progress of the

speech recognition. In this case, obviously utterance-by-utterance control won't do. The basic mechanism for handling this kind of movements in the Olga system, is the possibility to, at any specific moment, specify a parameter trajectory as a list of time-value pairs, to be evaluated immediately. Using such trajectory commands, gesture templates can be defined by grouping several commands together as procedures in a high-level scripting language (Tcl/Tk). This allows for complex gestures, such as "shake head and shrug" or "point at graphics display", that require many parameters to be updated in parallel, to be triggered by a simple procedure call.

4.4 Arguments for Pointing and Shrugging

Since a general scripting language is used, gesture templates can also be parameterised by supplying arguments to the procedure. For example, a pointing gesture might take optional arguments defining direction of pointing, duration of the movement, degree of effort etc. As another example, a template defining a "shrug" gesture can have a parameter for selecting one from a set of alternative realisations, ranging from a simple eyebrow movement to a complex gesture involving arms, head, eyebrows and mouth corners. During the course of the dialogue, appropriate gestures are invoked in accordance to messages sent from the dialogue manager. There is also an "idle loop", invoking various gestures when nothing else is happening in the system. The scripting approach makes it easy to experiment with new gestures and control schemes. This sort of template based handling of facial expressions and gestures has proven to be a simple, yet quite powerful way of managing non-speech movements in the Olga system.

5. THE DIRECT MANIPULATION INTERFACE

All visualisation in the Olga system except for the Olga agent is controlled by the Direct Manipulation Interface, DMI. It manages graphics output as well as user initiated input. The output may be used for displaying interactive menus, information tables and/or visualisations, e.g. photos of specific microwave ovens. The graphics component of the DMI is based on the Distributed Interactive Virtual Environment developed at the Swedish Institute of Computer Science [14], which simplifies real-time manipulation of displayed 3D objects.

6. THE SPEECH RECOGNISER

The Olga project was originally planned for two years, where the addition of the speech recogniser was scheduled for the second year. The intention was to make Wizard-of-Oz simulations with the Olga system during the first year of the project in order to collect speech and language material to be used for the training of the recogniser. However, due to various circumstances it later became evident that an Olga demonstrator had to be built during the first year. In order to get a better conception of Olga's intended functionality, it was decided to include a preliminary speech recognition facility.

The speech input module is based on the Waxholm recogniser described in Ström [15]. This is a software only continuous speech recognition engine with different modes for the phonetic pattern matching. In particular, standard multiple Gaussian mixtures and artificial neural networks are implemented for phone probability estimation. Thus, recognition may be performed either in a standard HMM or in a hybrid ANN/HMM framework. A lexicon with multiple pronunciations and a class bigram-grammar is used. The lexicon and grammar constraints are represented by a lexical graph, optimised for efficient lexical decoding. The decoding is performed in a two-pass search. The first pass is a Viterbi beam-search and the second is an A* stack-decoding search. Multiple recognition hypotheses can be output either as standard N-best lists or in a more compact word-graph format.

The recogniser was modified to interact with the dialogue interaction manager and speech input was enabled over the Internet. The current version of the Olga speech recogniser is very preliminary and only able to recognise sentences according to the written scenario that forms the basis of the Olga demonstrator.

7. ACKNOWLEDGEMENT

The original concept of Olga came from Henrik Wahlforss. He initiated and managed the project through the company Nordvis AB. The following organisations also took part: Centre for user oriented IT Design and Department of Speech, Music and Hearing, both at KTH; Department of Linguistics at Stockholm University and Swedish Institute of Computer Science. We would especially like to thank Eva-Marie Wadman and Olle Sundblad for graphical interface design and implementation. The Olga project was funded by Telia Research AB, NUTEK and Stiftelsen för kunskaps- och kompetensutveckling from September 1995 until September 1996.

8. REFERENCES

- [1] J. Beskow, M. Dahlquist, B. Granström, M. Lundeberg, K-E. Spens, & T. Öhman, "The Teleface project - Multimodal speech communication for the hearing impaired", in these *Proceedings of Eurospeech '97*, Rhodes, Greece, 1997.
- [2] J. Cassel, M. Steedman, N. Badler, C. Pelachaud, M. Stone, B. Douville, S. Prevost, B. Achorn, "Modeling the Interaction between Speech and Gesture", *Proceedings of 16th Annual Conference of the Cognitive Science Society*, Georgia Institute of Technology, Atlanta, USA, 1994.
- [3] N. Katashi, and T. Akikazu, "Speech Dialogue with Facial Displays: Multimodal Human-Computer Conversation", *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics (ACL-94)*, pp. 102-109, 1994.
- [4] K. R. Thórisson, "Gandalf: An Embodied Humanoid Capable of Real-Time Multimodal Dialogue with People", *Proceedings of First ACM International Conference on Autonomous Agents*, Marriott Hotel, Marina del Rey, California, pp. 536-7, February 5-8, 1997.
- [5] J. Bertenstam, J. Beskow, M. Blomberg, R. Carlson, K. Elenius, B. Granström, J. Gustafson, S. Hunnicutt, J. Högborg, R. Lindell, L. Neovius, A. de Serpa-Leitao, L. Nord, and N. Ström, "The Waxholm system - a progress report", *Proceedings of Spoken Dialogue Systems*, Vigsoe, Denmark, 1995.
- [6] S. McGlashan, "Towards Multimodal Dialogue Management", *Proceedings of Twente Workshop on Language Technology 11*, Enschede, The Netherlands, 1996.
- [7] W. Eckert, and S. McGlashan, "Managing Spoken Dialogues for Information Services", *Proceedings of Eurospeech '93*, Berlin, Germany, vol. 3, pp. 1653-1656, Sept. 21-23, 1993.
- [8] E. Giachin and S. McGlashan, "Spoken Language Dialogue Systems", in G. Bloothoof and S. Young (eds) *Corpus-based Methods in Language Processing*, Kluwer, The Netherlands, 1997.
- [9] J. Beskow and S. McGlashan "Olga - a conversational agent with gestures", to appear in *Proceedings of Workshop on Animated Interface Agents: Making them Intelligent*, Nagoya, Japan, August 23-25, 1997.
- [10] B. Le Goff, T. Guiard-Marigny, M. Cohen & C. Benoit, "Real-time analysis-synthesis and intelligibility of talking faces", *Proceedings of the Second ESCA/IEEE Workshop on Speech Synthesis*, New Paltz, New York, USA, Sept. 1994.
- [11] R. Carlson and B. Granström, "Speech Synthesis", in Hardcastle W. and Laver J. (eds) *The Handbook of Phonetic Sciences*, pp. 768-788, Oxford, Blackwell Publishers Ltd, 1997.
- [12] F. I. Parke, "Parametrized models for facial animation", *IEEE Computer Graphics*, vol. 2, nr 9, pp. 61-68, 1982.
- [13] J. Beskow, "Rule-based Visual Speech Synthesis", *Proceedings of Eurospeech '95*, Madrid, Spain, vol. 1, pp. 299-302, Sept. 18-21, 1995.
- [14] O. Hagsand, "Interactive MultiUser VEs in the DIVE System", *IEEE Multimedia Magazine*, vol. 3, nr 1, 1996.
- [15] N. Ström, "Continuous Speech Recognition in the WAXHOLM Dialogue System", *TMH-QPSR*, Department of Speech, Music and Hearing, KTH, Stockholm, Sweden, vol. 4, pp. 67-95, 1996.